

TITLE

Dynamic Time Division Multiplexing of Voice Signals Into Ethernet Data Streams Over A WAN Circuit

CLAIM OF PRIORITY

This application claims priority of and herein incorporates by reference in its entirety U.S. provisional patent application 60/262,722 filed on January 19, 2001 entitled "Dynamic Time Division Multiplexing of Voice Signals Into Ethernet Data Streams Over A WAN Circuit."

FIELD OF THE INVENTION

The present invention generally relates to transmission of voice signals by way of Ethernet data streams over a WAN circuit.

BACKGROUND OF THE INVENTION

The present invention solves problems relating to transmission of Ethernet Frames and Voice Signals (Voice, Fax, Modem – hereinafter called "Voice" Signals) over a WAN Circuit in a consistent manner (i.e. no Jitter) for Voice Signal Packets and therefore, preserves the quality of these signals while at the same time allowing for the maximum effective use of circuit bandwidth for the transmission of Ethernet Frames.

Ethernet Frames vary in size and can be over 1518 bytes in length. Frames of this size can monopolize a WAN Circuit causing degradation of the quality of Voice Signals sharing this circuit. A mechanism is therefore needed that allows for the maximum effective use of the bandwidth available in the WAN Circuit for both Voice and Ethernet Packets and which would ensure that Voice Signals never wait too long for the circuit. This mechanism should also operate with all standard Ethernet Frame sizes in a manner transparent to the attached Ethernet based communications devices.

Prior art systems and methods have employed the Local Subscriber Loop of the existing telecommunications infrastructure. Four types of prior art solutions exist that allow for the mixing of Voice Signals and Ethernet Signals on this circuit.

5

One prior art solution is to separate these signals within the spectrum available on the Local Subscriber Loop. This technique is used by systems which use frequency splitters (filters) to allow for low frequency analog Voice signals to co-exist on the Local Subscriber Loop with higher frequency digital data signals. This solution is widely deployed in the form of ADSL equipment with "POTS Splitters" (filters) attached at each end of the Local Subscriber Loop to enforce the separation of these frequency ranges.

10

1078713.1

The method used by ADSL equipment requires the addition of a filter circuit or POTS Splitter at each end of the Local Subscriber Loop. This separation of frequency bands on the media is not possible on other WAN circuit types such as SDSL, T1, and E1. In addition, the need to separate the spectrum of the analog Voice signals and the digital Ethernet signals on the same circuit requires partitioning of the spectrum in an asymmetrical manner. Digital signals coming downstream to the CPE device are provided at a faster rate than the digital signals going upstream from the CPE device. This asymmetrical data service creates problems for users who want to transmit large files upstream, as the speed of the data transfer in that direction is significantly slower than the downstream speed.

20

A prior art second solution is to convert Voice and Ethernet packets to ATM Cells for transmission over the Local Subscriber Loop. Digitized Voice Packets and Ethernet Packets are segmented into 53 byte ATM Cells. These Cells share the Local Subscriber Loop and are transmitted using xDSL technology. There is usually a priority scheme that provides that Voice Packets traverse the circuit with acceptable delays. This method requires ATM switches, ATM backbones, and ATM Voice Gateways to complete the solution.

25

ATM developers compromised between the needs of LAN data devices and Voice equipment and defined circuits based on 53 byte cells (48 bytes of payload) and then developed solutions that would fragment Ethernet packets to fit into them. This payload size adds jitter to the Voice data that circuitry must correct. Their efforts included transmission of signals through large systems of equipment. This meant that they had to develop methods and protocols which would establish a circuit and with required characteristics these systems. These protocols add complexity to this solution. Developers approaching this problem from a traditional telephony perspective had available to them circuits with fixed channels and developed approaches that would fit Ethernet packets into these fixed channels.

The method used by ATM cell based equipment creates a significant overhead or "Cell Tax" on each Voice or Ethernet Packet sent over the Local Subscriber Loop. The size of these cells also produces jitter in the transmission of the Voice data that must be smoothed out by adding circuitry to the solution. In addition to the reduction in available circuit bandwidth created by the ATM overhead, the systems that use this approach also require complex and expensive ATM system elements throughout. ATM is costly to develop and to deploy and is complex to install, configure and maintain.

A third prior art solution is used over T1 or E1 WAN circuits. This solution requires that the circuit be channelized. The circuit is divided into fixed size Time Division Multiplexed channels. Each channel provides a virtual connection of 64 Kbps over the WAN circuit. Devices using this technique fragment the Ethernet Frame into packets that fit into these channel sizes. Voice packets are sent in the same size channels as Ethernet packets. The bandwidth of this circuit is apportioned by assigning these channels to Voice or Ethernet packets.

The T1/E1 "Drop and Insert" method requires a channelized WAN circuit. There is significant additional cost in terms of circuitry and packet processing overhead that makes this a costly solution. Due to its fixed channel assignment nature, the bandwidth of the WAN circuit may not be used effectively to support the mix of Voice and Ethernet packets required by the

application. The channel assignments are fixed, and so go unused when there are no packets of the correct type to send.

A fourth prior art solution is called Voice Over IP (VOIP). This creates a type of IP packet to carry Voice data. These packets use all standard TCP/IP mechanisms to move through the network, including the WAN circuit. In addition, specific new protocol features are added to provide the "Quality of Service" (QOS) necessary for these Voice packets to traverse the network fast enough to maintain the quality of Voice signals.

VOIP developers are defining a system solution much like the ATM developers defined which would address large systems of equipment carrying mixed Ethernet Frames and Voice data. However, it does not focus on the Local Subscriber Loop. Additionally, VOIP requires that every element of the network infrastructure implement the same set of QOS features in order to produce reliable Voice support. The work to accomplish this through establishment of Standards and implementation of these Standards is underway but will require a long time and large expenditures to realize.

The present invention solves these and other problems in that it works with all existing Ethernet equipment as well as all existing analog Voice equipment requiring no changes in either. It can be deployed immediately and easily without the large expenditures necessary to modify or replace major elements of the supporting communications infrastructure.

SUMMARY OF THE INVENTION

The present invention injects and/or removes Voice data into and/or out-of an Ethernet stream in a dynamic, Time Division Multiplexer-like manner thus providing optimum quality Voice with minimal impact on Data bandwidth.

It is an object of the present invention to allow many services to be provided over the same network wherein the services are be easy to subscribe to and reconfigure under the control of the user of the services.

5 It is a further object of the present invention to supply a system and method that preserves the simplicity of Ethernet and adds to it a way to carry Voice (Telephony) signals that provides a smooth arrival rate, and hence high quality.

BRIEF DESCRIPTION OF THE DRAWINGS

10 Figures 1-8 depict various aspects and features of the present invention in accordance with the teachings herein.

DETAILED DESCRIPTION

The aspects, features and advantages of the present invention will become better understood with regard to the following description with reference to the accompanying drawings. What follows are a preferred embodiments of the present invention. It should be apparent to those skilled in the art that the foregoing is illustrative only and not limiting, having
20 been presented by way of example only. All the features disclosed in this description may be replaced by alternative features serving the same purpose, and equivalents or similar purpose, unless expressly stated otherwise. Therefore, numerous other embodiments of the modifications thereof are contemplated as falling within the scope of the present invention as defined herein and equivalents thereto. Use of absolute terms, such as “will not,” “will,” “shall,” “shall not,”
25 “must,” and “must not,” are not meant to limit the present invention as the embodiments disclosed herein are merely exemplary.

Figure 1 is a representation of major functional elements of a system that allows Ethernet LAN equipment and Voice equipment to share a single WAN circuit from a Subscriber Premise (home, office, campus building, business, etc.) to a Provider Premise (Central Office, ISP POP, MIS Department) in a manner which provides signal quality for the Voice equipment and at the same time makes effective use of the WAN circuit bandwidth to support Ethernet LAN traffic.

Item (1) provides the Ethernet physical layer interface and media access controller (MAC) that connects the Subscriber Premise device to the Ethernet LAN. This is a Standard interface compatible with any and all Ethernet networking equipment.

Item (2) may be activated when there are signals from the Voice equipment. If no Voice signals are present, Ethernet Frames are transmitted through the equipment unmodified. Item (2) is a software algorithm that is responsible for fragmenting Ethernet Frames into labeled packets as it passes through the circuit. At the receiving end of the WAN circuit this algorithm uses the labels to reassemble the Ethernet Frame into its original size. A unique aspect of this algorithm is that the packet sizes that the Ethernet Frames are fragmented into are selected by the equipment according to the line rate in use on the WAN circuit. These values are contained in a table of values within the software. The values in this table are computed to ensure that the maximum time that an Ethernet packet can monopolize the WAN circuit is short enough to ensure that Voice packets are delivered smoothly through the system. For example, a CODEC operating at the nominal sampling rate of 64 Kbps produces a new 16-byte value each 2 milliseconds. To allow these 16 bytes to flow through the system uninterrupted, the algorithm takes into account the WAN line rate in use and the 16 bytes of Voice data and combines this information with the known overhead of the HDLC Controller. The formula is:

Packet Size in Bytes=[(.002 Seconds*WAN Line Rate in Bits per Second)/8 Bits per Byte]-26

For example,

Ethernet Packet HDLC Overhead =	2 Flag Bytes
	2 CRC Bytes
Algorithm Overhead =	1 Identifier Byte

Voice Packet HDLC Overhead =	2 Flag Bytes*
	2 CRC Bytes*
Algorithm Overhead =	1 Identifier Byte*
Voice Packet Data =	16 Bytes*
Timing Accuracy Factor =	26 Bytes

To support more than one Voice channel, 16 bytes plus HDLC Overhead and the Identifier Byte are added to the Timing Accuracy Factor for each new Voice channel desired.

For example: If the WAN Line Rate = 2,320 Kbps
Then Packet Size = 554 Bytes

This size of Ethernet packet combined with the algorithm in Item (6) ensures that the Voice data arrival rate is 64 Kbps regardless of the WAN line speed or LAN traffic.

The table of values this generates for a particular piece of WAN equipment varies according to the WAN line rates in use. Provided below are some examples for the SDSL based product of the present invention.

<u>Line Rate</u>	<u>Packet Size</u>
2,320 Kbps	554 Bytes
2,064 Kbps	490 Bytes
1,552 Kbps	362 Bytes
1,040 Kbps	234 Bytes
784 Kbps	170 Bytes
528 Kbps	106 Bytes
400 Kbps	74 Bytes
272 Kbps	42 Bytes

The technique for low overhead fragmentation and reassembly of Ethernet Frames is part of the operation of Item (2). Prior art systems and methods are very time consuming and costly in terms of the amount of memory a system requires because data is moved several times as it passes through a device. The fragmentation and reassembly scheme used in Item (2) of the present invention allows for Ethernet Frames to be written into memory only once on their way through the device, while at the same time these Frames are fragmented into the required packet sizes for the WAN circuit. When the fragmented packets are received at the other end of the WAN circuit, they need to be reassembled into the original Ethernet Frames for delivery to the

LAN. This requires a method for the receiving end device to know when it has a complete Frame. The software in Item (2) also implements a method for identifying a complete Frame.

Ethernet Frames are received by Item (1) and placed into memory in the device. The presence of the Frame in memory triggers the software in Item (2) to begin to send it to the WAN Transceiver (7) for transmission over the WAN circuit (8). Item (2) uses the Packet Size indicated by the WAN line rate and computes the location of the last byte in a packet of this size within the memory range of the buffered Frame and adds one to it. It Reads and holds this byte from the Frame and overwrites an Identifier into this byte in the Frame. Exemplary Identifiers may include:

- 01H Complete Ethernet Frame
- 02H First Packet of an Ethernet Frame
- 03H Middle Packets of an Ethernet Frame
- 04H Last Packet of an Ethernet Frame

The algorithm then passes the address and packet size information for this packet (which now has an additional byte added by the Identifier code) to the DMA controller. The DMA controller then handles the transfer of this packet from device memory to the WAN Transceiver that transmits it over the WAN Circuit. The software algorithm in Item (2) then writes the saved byte that was replaced by the Identifier back into its original location and this becomes the first byte of the next packet. This pointer manipulation technique fragments the Frame into the desired packet sizes without requiring multiple moves of the Frame in device memory.

On the WAN Circuit receive side, the packets are received by the Transceiver (9) and written into memory by the DMA controller. The software algorithm (11) then uses the Identifiers to recognize the type of packet and to determine when it has a complete Ethernet packet. As it reads the Identifiers they are removed and the original Frame is sent to the Ethernet LAN by the MAC and physical interface circuitry in Item (12).

This scheme supports Full Duplex operation, so descriptions of Items (1) and (12), (2) and (11), and (7) and (9) may be identical.

Item (3) provides a standard analog telephone line interface style FXS which allows for connection of any standard analog telephone, modem, or Fax equipment. Item (3) contains the physical layer interface and the Subscriber Line Interface Circuit (SLIC) required by Voice equipment. It also contains the CODEC that digitizes these analog signals for transmission through the device. The software in Item (4) controls the SLIC/CODEC circuits and handles Off Hook/On Hood and Ringing signaling. In addition this software transfers 16 byte Voice packets between Item (3) and Item (7), the WAN Transceiver. Voice packet Identifiers are added as the last byte of each packet and the restoration of the original packet stream at the receiving end removes these Identifiers. In the described embodiment, the Voice Packet Identifiers are 80H to 8FH. Bits 0-3 of this byte contain signaling information for the telephone line operation.

Item (5) supports two small Management tasks. As a test of WAN Circuit quality, a 32 byte Management packet is sent across this circuit every 1 seconds. If the receiving end gets the packet and can decode it, the Link LED on our devices is blinked as a visual indication that not only is the WAN Circuit up, but that it has sufficient quality to allow for the error free transmission of data. The function of Item (5), which pertains to mixing Ethernet and Voice packets over the WAN circuit, is that this 32-byte packet contains a code, which indicates that the equipment sending the packet is a unit that supports mixed Voice and Ethernet. This is done for compatibility. Optionally, if this code is absent, the devices will revert to Ethernet-only operation and none of the features which support mixing in Voice packets will be activated. This preserves all WAN Circuit bandwidth for the LAN traffic.

Item (6) controls the flow of all packets to/from the WAN Transceiver. It uses a priority scheme that provides high quality Voice signals. The scheme is very simple. If there are Voice packets available, they are sent. If there are no Voice packets available it sends one Ethernet packet and then looks for Voice packets again. The priority scheme is Voice 1st, Management 2nd, Ethernet 3rd. Management packets are generally sent on a 1 second interval.

Item (7) is the WAN transceiver that provides the physical and media access controllers that support the WAN circuit. The physical layer can be, for example, xDSL, clear channel T1, clear channel E1, Wireless, or FTTH (Fiber to the Home). Other known physical layers may be used as well. The media access controller may be an HDLC controller. This allows for reliable communications over the WAN circuit with the minimum necessary overhead, preserving more bandwidth for actual user traffic.

On the Provider Premises side of Figure 1 all the functions are duplicates of those shown on the Subscriber Premises side except with respect to Item (14). Item (14) can either convert the Voice packets back to analog Voice signals of the FXO type for connection to a telephone switch, or can keep the signals digital and send them into the PSTN infrastructure in T1 DS0's using standard T1 signaling techniques to control the telephone circuit.

In Figure 2, Items (16) and (17) graphically represents the mixing of Ethernet and Voice packets on the WAN circuit including details about the Identifiers used in each packet type. This diagram shows how the solution switches between traditional bridging operation where Ethernet Frames of any size that arrive are forwarded through the device and out the WAN circuit to our technique for dynamic time division multiplexing which ensures a 64 Kbps channel across the WAN circuit for Voice data for each Voice channel whenever Voice signals are active.

The solution of the present invention differs from that used in the prior art to support Ethernet and Voice using ADSL over Local Subscriber Loops in that it supports symmetrical upstream/downstream communications and will work for circuit types where dividing up the frequency bands is not possible such as SDSL, T1 and E1. The present invention is not limited by requiring that the spectrum of the circuit be divided between analog Voice signals and Digital Data signals. This difference also means that our solution does not require the "POTS Splitter", the filter necessary to ensure that the frequency bands are separated.

The present invention differs from that used for ATM based solutions in that the size of Ethernet packets sent across the WAN circuit can be significantly larger, requiring less complex

circuitry and reducing the overhead which reduces the available bandwidth on the WAN circuit. It also delivers a smooth stream of Voice data that reduces the complexity of the circuits necessary for handling the Voice data. ATM also has complex configuration requirements that are minimized in our solution. This allows operational connections to be established by simply connecting the wires. This is significant at the system and enterprise level as it reduces the costs to install, configure and maintain our solution.

This solution differs from that used in T1/E1 Drop & Insert equipment as it works over an unstructured (Clear Channel) WAN circuit and only reserves bandwidth for the Voice data when the Voice channel is actually in use. These circuits are typically much less costly to provision and hence to purchase as they require less circuitry to implement. In addition, our solution allows for much larger Ethernet packets to be sent across the WAN circuit resulting in greater effective bandwidth available for user data with less overhead lost to framing and lower system costs as less packet handling is required.

VOIP solutions encapsulate the Voice data into IP packets and then add additional protocols which use information contained in fields of the IP packets to signal to equipment that these Voice packets must receive priority treatment in order to maintain the arrival rate that Voice signals demand. This system perspective adds processing complexity to the solution that we do not think is necessary to solve the problem of mixing Voice and Ethernet data on a WAN circuit. Like ATM, the packet sizes of the Voice packets will be larger than the 16 bytes of our solution and so will require smoothing to remove jitter.

As shown in Figure 3, an exemplary Voice Over xDSL system architecture, the customer premises are connected to the Service Provider Premises via subscriber loops. The customer premises may include an Ethernet connection and telephone in communication with a CPE device associated with subscriber loops that may be owned or leased by a service provider. The subscriber loops are further connected to the service provider via a digital subscriber loop access. At the service provider premises a T1 circuit connects Voice from each access to Voice Switches.

Figure 4 is provided as an exemplary Voice Over xDSL Dynamic Time Division Multiplexed data packet in accordance with the present invention.

Figure 5 is provides exemplary Voice Over xDSL functional circuit elements. Voice is given priority and injected into a data stream using the dynamic TDM of the present invention. In this example, SLIC may provide 2-to-4-wire conversion; analog pulse shaping; off hook/on hook detection; and ringing voltage. CODEC may provide D/A conversion and off hook/on hook/ringing signals. T1 framer may provide extended super frame or super frame. The CPU voice function may provide off hook signal, on hook signal ringing signal as well as the ability to pass through all other signals.

Figure 6 is provided as an exemplary IMUX packet handling application wherein the IMUX can be used with any WAN interface; for example, IDSL, SDSL, T1 or E1. As shown, inverse multiplexing creates a single virtual physical layer interface. The IMUX packet handling logically groups two or more actual physical interfaces, creates WAN pipes 2x, 3x or even 24x the bandwidth of a single port and is used as if it were a single physical pipe. Packets are striped across the available bandwidth using a first available circuit algorithm. The IMUX packet handling ensures in order packet arrival as required by some applications. Additionally, if a circuit fails, the algorithm of the present invention continues to use all operational circuits and failed circuits are automatically added back in when they become operational again. Furthermore, by providing aggregated bandwidth little or no management through use of individual links is required.

Figure 7 is provided as an exemplary Voice Over xDSL TDM and IMUX application. As shown in Figure 7, data, voice and management packets share the xDSL circuits. The packets are striped across the xDSL circuits and the IMUX provides for in order arrival of packets.

Figure 8 is provided as an exemplary Voice Over xDSL with IMUX detailed architecture.

5

Having now described a preferred embodiment of the invention, it should be apparent to those skilled in the art that the foregoing is illustrative only and not limiting, having been presented by way of example only. All the features disclosed in this specification (including any accompanying claims, abstract, and drawings) may be replaced by alternative features serving the same purpose, and equivalents or similar purpose, unless expressly stated otherwise. Therefore, numerous other embodiments of the modifications thereof are contemplated as falling within the scope of the present invention as defined by the appended claims and equivalents thereto. For example, the techniques may be implemented in hardware or software, or a combination of the two as shown or in other equivalent embodiments.

It will be appreciated that the system, method of operation and product may vary as to the details of its configuration and operation without departing from the basic concepts disclosed herein.